

**Inside Intel's Itanium:
A Strategic Planning
Discussion**

An Executive White Paper

March 2000

Inside Intel's Itanium: A Strategic Planning Discussion

Executive Summary

More than a year before its scheduled arrival in the latter half of 2000, Itanium — Intel's first IA-64 microprocessor — is driving the long-range hardware strategic plans of many Information Systems (IS) executives for the new millennium.

And it should.

In this *Executive White Paper*, Aberdeen provides a general overview of Intel's IA-64 program and a detailed discussion of Itanium — the first IA-64 microprocessor available to IS executives. The paper is divided into the following sections:

- The significance of 64-bit computing and the key attributes of Itanium that differentiate it from RISC-based architectures;
- Why Itanium is important as the first microprocessor of an Enterprise Architecture;
- How Itanium will perform with key applications that are migrated to take advantage of the architecture;
- When the industry should expect to see Itanium systems, and how these systems compare to those of other architectures; and
- How IS executives should consider migrating to IA-64.

Based on our findings, Aberdeen makes the following conclusions about IA-64 and Itanium:

- Applications such as business intelligence, Internet commerce, and Online Transaction Processing (OLTP) may perform better by taking advantage of IA-64 parallel-processing technology;
- Itanium solutions will feature enhanced reliability that will help boost system performance for the high end due to Intel's focus on the addition of platform-level features;
- Independent Software Vendors (ISVs) will likely accept IA-64 as a de facto standard and will migrate their 32-bit software to Intel's IA-64 64-bit platform as their applications mature to the point of requiring 64-bit processing;
- Because of the factors cited above, over time IA-64 will become the volume leader for enterprise applications if Intel continues on its current publicly stated schedule — within five to seven years of its debut. Consequently, Aberdeen recommends that IS executives identify an IA-64 adoption strategy;
- Given that enterprises need to build up an IA-64 learning curve, Aberdeen suggests that IS executives deploy early platforms with Itanium in order to determine how their key enterprise applications will

perform on the new platform — and how to tune them to best fit the application; and

- Initial Itanium deployments will likely be best accomplished with *new* applications that will fully take advantage of Itanium's capabilities. Existing applications running on other CPU architectures should be planned for later migrations to IA-64.

Planning Considerations for IA-64

When delivered, Itanium-based systems should, in the short run, immediately boost the performance and scalability of today's high-end, Intel-based 32-bit server family. Itanium is the first of a family of processors, called IA-64, that will enhance the Intel architecture with 64-bit capabilities, build new performance enhancement features into the architecture, and maintain 32-bit compatibility in the hardware platform. However, it is in the long run that Itanium and follow-on IA-64-based processors should have their largest impact on IS, due to the following:

- Advances in chip technology, known as Explicitly Parallel Instruction Computing (EPIC). These advances will give IA-64 clear differentiation from RISC-based systems by more effectively using parallelism in processing — allowing users to apply IA-64-based systems more effectively on specific high-end applications such as Business Intelligence. Furthermore, IA-64 ties the CPU closer to the program compilers, thus in all likelihood allowing better performance optimization for high-end applications. Other elements of the IA-64 architecture will permit better memory utilization by applications.
- Closer coordination among the microprocessor, the hardware, and packaged-application suppliers *during chip design*, resulting in better tuning of the system for very-high-end, business-critical applications.
- Enhanced reliability of the platform that will also help boost system performance and scalability, especially at the high end. And Intel aims to add more platform-level features to bolster server reliability at the chip set and server I/O levels.
- The likely acceptance of IA-64 microprocessor technologies as the reference platform and industry standard during the first decade of the new millennium by ISVs, provided Intel continues to build its support from software vendors in the coming years. ISVs will gradually migrate their 32-bit software to Intel's IA-64 platform in order to take advantage of the new EPIC architecture.
- With the delivery of IA-64, IS executives who have traditionally deployed RISC-based (reduced instruction set computing) servers to meet enterprise business-critical needs will have their first opportunity to stan-

standardize enterprise hardware platforms on one architecture for server environments.

- Initial Itanium deployments will likely be best done with *new* applications that will fully take advantage of Itanium's capabilities. Mission-critical existing applications should be planned for future migrations to IA-64.

What Does "64 Bits" Mean, and Why is IA-64 More Than Just a 64-bit Chip?

For IS executives, the question is not why but when they should move enterprise applications to 64-bit platforms. Today's 32-bit computer architectures typically address 4 GB of main memory (especially Intel servers) and thus satisfy the performance requirements of a majority of today's applications. However, as Internet commerce and large database applications continue to grow in size, memory addressability requirements will increase in server and workstation applications. In response, 64-bit architectures increase by several orders of magnitude a system's ability to address main memory. The difference in performance can be seen in applications that require more than 4 GB of memory — for example, a number of large database, Internet ERP, and new line of business (LOB) applications. As memory prices drop, it will soon be economically possible to have main memory data storage in the tens of gigabytes. Combining this increased storage with a 64-bit address range will allow significant performance improvements for large databases, ERP, and future Internet applications.

IA-64 is not the first architecture to use 64-bit addressability. Current RISC architectures, including HP's PA-RISC, Sun UltraSPARC, Compaq Alpha, and IBM PowerPC, are already at 64 bits for key applications. But "64-bitness" is just one step. Platforms will also need to support very large memory (VLM) — larger than 4 GB — to take advantage of all of the benefits of 64-bit architectures. Intel is working with the industry to ensure that hardware, operating systems, and other software will take advantage of this platform capability.

IA-64: The Architecture

Beyond IA-64's VLM support, there are other traits that differentiate IA-64, including its new EPIC design philosophy that will handle parallel processing differently than previous architectures.

EPIC, which Intel developed jointly with Hewlett-Packard, changes how the microprocessor interacts with applications by feeding data more effectively into the processor. The design addresses a number of issues that other architectures, most of them designed 10 to 20 years ago, have not yet fully overcome, including achieving maximum parallelism in high-end situations and reducing microprocessor layout complexity to speed processor upgrades. IA-64 processors will also contain massive chip execution resources to support more of everything — regis-

ter, functional units, logic, instructions, cache, and data — translating to additional processing power well beyond what is available today on other processors. The appendix of this paper provides in-depth details of the IA-64 architecture.

Why Itanium Is Important

Itanium will be the first commercially available Intel 64-bit architecture, representing what's to come in the high-end server marketplace. The entry of Intel into other microprocessor markets has resulted in commoditization at the Central Processing Unit (CPU) level, as Intel's volume manufacturing methods have reduced CPU costs and improved performance. And these resulting changes in price and performance have created considerable positive momentum for mass-produced Intel microprocessors on a global scale. Intel's entry at the high end of the server market will result in downward pressure on high-end server prices and continued rapid price-performance decreases at the high end for enterprise IS executives.

But is Intel's Itanium architecture the right choice for early adopters? Most Information Systems (IS) strategic planners have already heard about Itanium's changed schedule, which moves availability from late 1999 to the second half of 2000 due to a slower-than-expected chip-design stage. Despite this delay, Intel states that its Itanium chip set is now on track for a mid-2000 introduction. Intel intends to follow up Itanium with additional IA-64-architecture microprocessors: McKinley in 2001, and Madison and Deerfield in 2002.

The answer to the question of when IS executives should deploy applications on the IA-64 architecture will likely depend on whether the application is an existing application, and whether it is to be migrated from another hardware platform (existing RISC architectures or IA-32 architectures). Itanium represents the earliest opportunity to deploy the IA-64 architecture. For those who are deploying new applications, and who wish to run those applications on IA-64 in the long run, Itanium will be a good place to start. By starting early, users can avoid migrating the application down the line.

Thus, for ISVs anxious to take advantage of the new functionality of IA-64's architecture, known as Explicitly Parallel Instruction Computing (EPIC), Itanium looks like the right platform-of-choice for early application development and prototype deployment. The focus of a following section, "How Itanium Will Perform on Key Applications" is how users can determine what new applications are leading candidates for Itanium deployment.

The Competitive Landscape for Itanium's Release

Aberdeen suggests that it is too early to clearly determine how Itanium will perform against other microprocessor architectures on the market at Itanium release time, although Intel estimates that Itanium will be more than competitive against

existing enterprise RISC architectures. However, it is clear that there will be a significant number of alternative architectures available. As a result, it will take time for IA-64 to gain momentum.

Among other competitors, IBM will continue to push the performance envelope with its PowerPC microprocessor architecture, making headlines with its Silicon-On-Insulator and copper CPU-technology innovations that will keep the S/390 and AS/400 platforms competitive. Sun is continuing to gain market share in the commercial application Unix server market (especially in Internet servers) with successive generations of its UltraSPARC architecture and more compact (and powerful) microprocessors. Hewlett-Packard will continue to offer its PA-RISC-based systems until such time as IA-64 can outperform PA-RISC-based systems. Compaq continues to design new generations of its Alpha microprocessor, and expects to deliver the EV7 (at roughly the same time as Itanium) that is projected to reach frequency speeds of over 1 GHz.

Despite its competition, IA-64 has seen support from all of the aforementioned system vendors, with the exception of Sun. In fact, several vendors have committed to discontinue their RISC processor lines in favor of IA-64. Aberdeen notes that IBM, Sun, HP, and Compaq all have announced plans to run their Unix operating environments on IA-64 — a factor that will contribute mightily to mass adoption of IA-64 in the marketplace.

Key Ingredients for an IA-64 System

During the development of its IA-64 components, Intel has worked closely with key industry suppliers to ensure that complete solutions are available when Merced ships. And IA-64's new architecture design, based on EPIC, means that IS executives will likely be offered complete hardware, applications, and operating system platforms from suppliers instead of separate components provided in piecemeal fashion. These ingredients, in addition to the Itanium-based server hardware, include operating systems, application development tools (compilers), and operating systems that have been "64-bit-ized." By the time that IA-64 is released, these components will have been recompiled to take better advantage of the architecture, while still offering IS executives backward compatibility to 32-bit applications (Figure 1).

IA-64 Industry Support

Aberdeen notes that IA-64 has received unprecedented industry support. Figure 1 clearly demonstrates this support, which includes:

- Almost all major high-end system vendors, except Sun Microsystems, committing to deliver IA-64 systems. Combinations of Intel Architecture vendors like Dell and Compaq, and traditional RISC vendors like HP, SGI, and IBM have all stated intentions to deliver Itanium-based systems.

A number of these companies, including HP and SGI, have committed to fully transition their product lines to IA-64 over time.

- Eight Operating System Vendors (OSVs), including Microsoft (Windows 2000), Novell (NetWare), and various suppliers of Unix, have stated their support for IA-64. A significant part of the porting is already complete, according to Intel. All are booting IA-64-enabled versions of their operating systems using Intel's pre-silicon software tools.
- Public ISV commitments from many of the largest server and workstation application vendors have also been offered in support of IA-64. Vendors include Oracle, SAP, PeopleSoft, Baan, Microsoft SQL Server, IBM DB2, and others. Intel states that some applications are also running on the pre-silicon environment.

Figure 1: IA-64 System Components

64-Bit-Enabled Applications	<ul style="list-style-type: none"> •Business Intelligence •The Internet •Databases •LOB •Scientific/Technical
Software/Compilers Application Development Tools	<ul style="list-style-type: none"> •Tool Vendors •Application Libraries
64-Bit-Enabled Operating System	<ul style="list-style-type: none"> •IBM/SCO Monterey •Microsoft Win64 •Novell Modesto •Sun Solaris •HP-UX, Linux, etc.
IA-64 Hardware EPIC IA-32 Binary Compatibility	Compaq, HP, IBM, Dell, NCR, SGI, Unisys, Sequent, Bull, Siemens, NEC, Hitachi, Fujitsu

Source: Aberdeen Group, July 1999

System suppliers tell Aberdeen they will offer key packaged applications that take advantage of the larger 64-bit addressability within the IA-64 architecture, including large-database applications, the Internet, Business Intelligence, and technical/scientific applications (see "How Itanium Will Perform on Key Applications" section). Operating system suppliers are planning to provide versions of their operating systems to run on IA-64 (see "What IS Executives Need to Do to Migrate to IA-64" section).

The Broader View on ISV Development of IA-64 Solutions: The IA-64 Fund

Intel recognizes the need for complete application solutions for Itanium platforms. To foster application development by ISVs, especially in Internet and emerging enterprise spaces, Intel is co-investing with a number of system suppliers and select large Fortune 500 IS executives in an IA-64 Fund that will be used to help ISVs develop applications for Itanium and future IA-64 processors. Together, suppliers and IS executives will target a portfolio of specific ISVs, bridging the gap between technology development and IS executives who will deploy these application platforms with Itanium. The fund, valued at \$250 million, complements a number of other initiatives currently underway, including software-enabling programs, application solution centers (ASCs), Intel Itanium-based Software Development Kits (SDKs), and Enterprise Technology Centers.

Implications for IS Executives' Hardware Deployment Strategies

Aberdeen notes that with this level of industry support the breadth of packaged solutions available to IS executives will be ever increasing, as Itanium and future versions of the IA-64 architecture are rolled out by Intel and system suppliers. Aberdeen suggests that as more packaged solutions become available to IS executives, packaged-application support — combined with increasingly superior performance and price-performance — will drive IA-64 adoption. That is why it is important today for IS executives to identify an IA-64 adoption strategy. What remains to be clarified is whether this strategy should begin with Itanium, or if IS should wait for future versions of IA-64 to reach the marketplace.

Intelligent Hardware Reliability: Migrating System Robustness to Hardware

Enterprises today require exceptionally high robustness for business-critical production environments. Today, most large enterprises run business-critical applications on a combination of mainframe and RISC-based servers that support Unix environments, and Intel Architecture servers are starting to grow in importance as well. RISC-based server suppliers have been gradually introducing hardware reliability features to their systems, and over time Intel has enhanced its IA-32 architecture with hardware-based reliability features to reduce system downtime. IS executives have also used other tactics to achieve reliability, including the use of enterprise fault-tolerant servers as well as various forms of high-availability clustering.

IA-64's introduction continues a drive now underway toward better server reliability. To achieve what Aberdeen calls "intelligent hardware reliability," Itanium employs a number of tools built into the hardware to handle error detection, correction, and "containment." These features include enhanced Error Correction Code (ECC), an enhanced machine-check architecture, improved processor-firmware-OS integration, and a host of other improvements. (For more details on these features, see "Itanium Intelligent Hardware Reliability Features" in the appendix.)

The availability of enterprise class Unix on IA-64 will also enhance platform reliability. A number of high-end vendors (HP, IBM, Unisys, and others) are committed to validating IA-64 platforms for their high-end enterprise applications.

How Itanium Will Perform on Key Applications

As a 64-bit architecture, Itanium will be especially appropriate for new and high-end applications — those for which migration effort is nil, and those able to take full advantage of Itanium system features that include VLM (very large memory) and enhanced reliability features. Even within the 64-bit architecture space, Itanium's differentiation means that it will deliver performance and scalability beyond other architectures in certain key application areas, particularly the following:

- Business Intelligence;
- The Internet;
- New line-of-business (LOB) applications; and
- Technical/scientific applications (i.e., those supported by technical workstations).

Business Intelligence

To support terabyte-and-beyond data warehousing and data mining, systems need VLM support. They also need a high degree of parallelism to ensure that even complex transactions do not slow the bulk-load/process/bulk-store operations typical of large-scale and complex queries. Itanium not only provides 64-bit VLM support, but its predication techniques (described in the appendix) also allow exceptional parallelism even of transaction code with unpredictable control flows. Moreover, the Level 3 cache speeds CPU processing of large datasets by reducing time-to-load-from-memory for a larger amount of data versus disk I/O speeds.

The Internet

The Internet involves both multimedia (text, audio, video, graphics) and exceptionally rapid variation and increase in transactions (queries and OLTP). Thus, Internet systems need to scale rapidly for both decision support and e-commerce. As with Business Intelligence, Itanium's predication allows higher performance on large-scale queries or Java object code with unpredictable control flows. Itanium's speculative loads from (and stores to) main memory likewise enhance parallel performance beyond the norm. Examining Intel's preliminary performance results, Aberdeen concludes that Itanium can deliver performance and scalability even beyond today's 64-bit RISC systems in both Internet-style queries and OLTP.

Users should also note that Itanium provides specific features for Internet applications: security (e.g., encryption algorithm support) and streaming media (e.g.,

floating-point performance improvements and support for applying a single instruction to multiple larger scale data streams).

New LOB Applications

In keeping with Aberdeen's suggestion that a Itanium deployment should focus on new-application opportunities rather than the migration of existing applications, Itanium's architectural enhancements will benefit most new LOB applications. Many of these applications use small or predictably sized integer data. Itanium's instruction pipelining and large register set allow greater parallelism of this small-size data, increasing performance in many key LOB-application operations. Itanium's support for both queries and OLTP also aids LOB applications, which are often "mixed-workload." Itanium's reliability enhancements will benefit end-users that rely on these key applications.

Technical Workstations

Itanium's floating-point features, 64-bit addressability for large multimedia data-sets, and increased parallelism for the complex processes typical of technical workstations, all combine to deliver new performance and scalability benefits to a market already highly receptive to Intel architectures. The floating-point registers will provide better performance for graphics, as well as extended precision for performance on larger data sets and computations.

When IS Executives Should Expect to See Itanium-based Systems

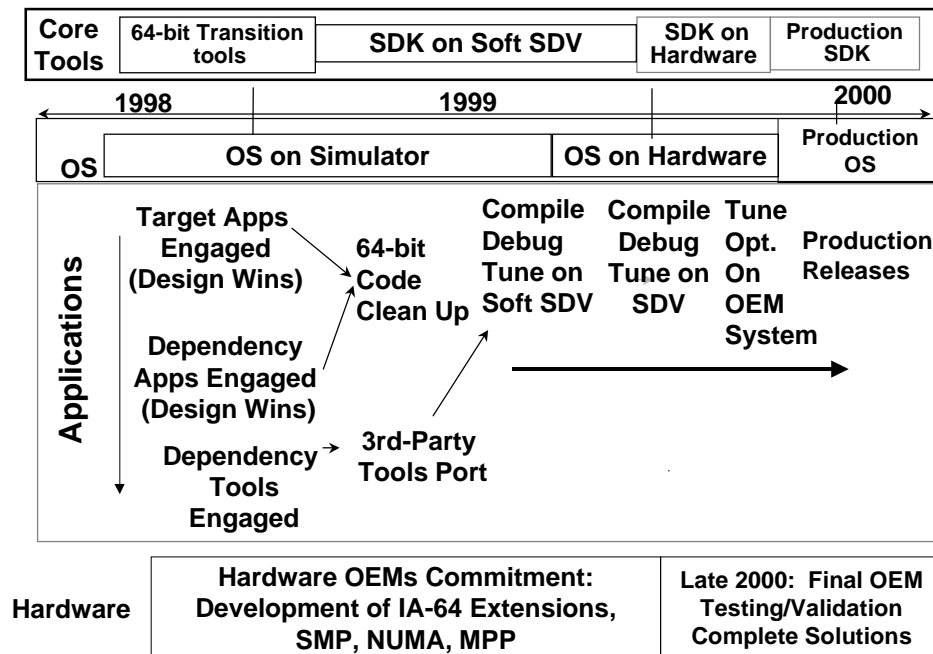
As previously stated, the road to Itanium is generating enormous industry momentum. Key components need to be put in place before IS executives will be able to fully utilize Itanium-based systems and workstations and servers. These components include applications that take advantage of IA-64's 64-bit addressability, operating systems that are 64-bit-enabled, compilers that have been optimized to use with EPIC, and hardware systems that support Itanium. Most are now in development, and today it is a race to the wire. Savvy IS executives should monitor Itanium's release process because, as with any industry initiative, Itanium's arrival date could change.

Today, Intel states that it will meet its "second-half-of-2000" deadline for Itanium's delivery. As of June 1999, hardware development efforts around Itanium are aligned — with both chip set and processor components on track. Furthermore, Itanium supplier OEMs have created more than 30 workstation and server designs that are now being finalized. And, as the following section will elaborate, multiple applications and operating systems are now being optimized for the Itanium platform.

During 1999, Intel expects to build a large pool of ISVs that have compiled their applications to run on Itanium and IA-64. After having worked with key ISVs, Intel

will provide 64-bit transition tools to the entire ISV community, to assist developers in 64-bit-enabling their applications. Key operating system suppliers are already running their 64-bit-enabled OS on Itanium's simulation platform (Figure 2).

Figure 2: IA-64/Itanium Release Roadmap



Source: Intel Corp., July 1999

What IS Executives Need to Do to Migrate to IA-64

While many of the initial Itanium and IA-64 deployments will be for new applications, Aberdeen realizes that different IS requirements will drive a variety of application migrations. To help identify the right migration strategy and time frame, Aberdeen recommends that all IS executives planning to migrate to IA-64 systems take the following steps to prepare:

1. Evaluate which in-house applications will benefit from IA-64 and what operating systems they run on — are both ready for IA-64? If not, when will they be ready?
2. If the customer has a choice between IA-64 system suppliers, evaluate which system will require less migration effort; which will allow the user to take more immediate advantage of IA-64's performance and reliability features; and which system supplier provides greater support for the migration effort. The key is to begin asking system providers today about their IA-64 plans — and identify those organizations with the most robust experience.

3. For key performance-critical applications, begin to test the effects of IA-64 using supplier-provided compilers. Users forging partnerships with systems suppliers may want to arrange to take advantage of these suppliers' own testing environments.
4. Create a plan for migration, including specification of testing systems and a time line for moving migrated applications into production. This plan should involve running most applications initially on a 32-bit compatibility box, and rewriting key performance-critical applications to take advantage of IA-64.

Once enterprises have identified targets for migrations, IS executives should also evaluate for planning purposes when it makes sense to migrate to IA-64, especially if Intel's plans change again. For most IS executives, migration will involve a recompile for most internal applications and an upgrade from the ISV for packaged applications. The good news for packaged applications is that most ISVs are being exceptionally proactive in delivering versions of their products supporting Itanium, but the customer should check with the ISV anyway. Migration efforts will vary by the type of system presently used:

- Windows NT;
- Intel 32-bit Unix; or
- RISC 64-bit Unix.

Windows NT

Microsoft has pledged to provide 64-bit Windows 2000 when Itanium arrives. Intel states that 64-bit Windows 2000 will allow existing IA-32 applications to run as they are on Itanium and IA-64 systems. Microsoft also indicates that most, if not all, applications written to IA-32 should require only a recompile to run on IA-64. Thus, migrating a Windows-NT-based application should require a recompiled version of the application and a system upgrade to Windows 2000 64-bit operating system. Aberdeen recommends that mission-critical Windows NT applications be recompiled (typically to ensure that performance-critical variables can run at 64-bit and the address space can be greater than 4 GB for maximum performance).

Intel Unix 32-bit

With the help of IBM, suppliers such as SCO and Sequent are making extensive efforts under the heading of "Project Monterey" to support their IS executives' migration efforts. IBM is working with both suppliers to combine AIX with SCO's UnixWare and Sequent's DYNix. Users can run IA-32-compliant applications (without needing a recompile) on an IA-64 box, and suppliers are also offering two target programming models: 32-bit on 64-bit (called ILP32), and full 64-bit (called LP64). The 32-bit on 64-bit model will in most cases require a recompile of

the application to run on IA-64; the full 64-bit model will also require rewriting of legacy applications that are too specific about variable sizes. Aberdeen recommends that most applications be recompiled using 32-bit on 64-bit, and that performance-critical applications requiring large memory space use full 64-bit.

RISC Unix 64-bit

Suppliers such as IBM (AIX), HP (HP-UX), and Sun (Solaris — through indirect channels), aim to offer Unix operating systems on Itanium and to provide extensive services for users of alternative 64-bit architectures looking to migrate to IA-64. These suppliers also provide specific tools to migrate existing applications on 32-bit systems to IA-64. For IS executives, migration may involve moving to the supplier's IA-64 Unix as well as from 32-bit to 64-bit systems. For example, Sun IS executives may want to move to 32-bit Solaris on Intel now, then Solaris Itanium at a later date. Information Technology (IT) buyers with existing 64-bit systems should be able to run their applications on IA-64 without needing to recompile. In many of these cases, applications will not require rewriting. In either case, Intel suggests acquiring recompiled applications or recompiling in-house code for optimal performance.

How Itanium Addresses 32-bit/64-bit Compatibility

Itanium represents a significant shift for the Intel Architecture, much like the shift the industry has taken in moving from 16-bit to 32-bit platforms — a process that has stretched close to 10 years (and vestiges of 16-bit desktop applications continue to exist today). Because of the magnitude of the transition, one of Itanium's key aims is to facilitate that transition while giving IS executives a taste for what IA-64 can accomplish. Future IA-64 microprocessors will build on Itanium's foundation by adding new layers of performance enhancements.

IA-64's architects have designed Itanium to provide backward binary compatibility to IA-32 applications that run on Intel's X86 platform. A portion of Itanium's hardware real estate is dedicated to processing 32-bit code, with portions of the CPU communicating with the compilers to shift gears to process the 64-bit code.

IA-64's IA-32 compatibility is meant solely for compatibility, not necessarily maximum performance. IS executives looking for maximum performance for an application will require recompiled IA-64 versions of the operating system and applications. For maximum IA-32 performance, IS executives should strongly consider Intel's IA-32 architecture instead of IA-64.

Aberdeen Conclusions

The need for more computing power is spiraling upward as data-intensive decision support, Internet-based multimedia, and transaction-oriented applications gobble up as much computing power as an enterprise can put forward. Intel's IA-64 architecture is expected to help Information Systems executives keep pace

with user demands for more power and better response times due to new IA-64 technological advances such as Explicitly Parallel Instruction Computing (EPIC).

But Aberdeen suggests that hot technology is not enough to drive the success of IA-64. IA-64's industry support is likely to be the key factor in its success. The OEM and OSV support has been unprecedented. Intel and its partners such as Hewlett-Packard, IBM, Compaq, and many others have all been working aggressively with their ISVs to ensure that complete application/OS/hardware platforms are available for Itanium system production. And, by so doing, the systems vendors and Intel may be taking much of the transitional effort of moving from 32-bit to 64-bit computing off the shoulders of the IS department.

Itanium, the first IA-64 processor, should be viewed as a strategic move for IT departments. It promises performance enhancements with applications like OLTP, decision support, and technical computing environments like CAD, DCC, etc. Aberdeen suggests that Itanium will excel in new application environments, including e-business and large in-memory databases. First, these applications are likely to be the easiest to deploy because they do not require migration. Second, these new applications will give IT the opportunity to evaluate IA-64's potential for broader deployments in their enterprises.

If Intel's IA-64 captures the industry's leading ISV applications, the high-end of the computing market could experience the same "commoditization" effect at the microprocessor level (and then systems level) that the industry saw when Intel moved into the PC marketplace and the low-end/midrange server marketplace. Volume manufacturing tends to drive costs down — which could signal some real problems for chip-makers that wish to compete with Intel for high-end business over time. Still, commoditization has resulted in an overall lowering of computer hardware costs as well as increased competition on the performance side of the equation — and the lower costs/higher performance dynamic has been great for the IS buyer.

Challenges for Itanium remain. The greatest challenge is meeting Itanium's delivery date — fulfilling a schedule that requires consensus among the application suppliers, operating system suppliers, and system suppliers to build a monolithic, strikingly different enterprise platform. The next challenge is making sure that IA-64 EPIC compilers are optimized to handle key applications — and work closely with the microprocessor to provide the compelling performance that IS executives demand in enterprise environments. A third and final hurdle will be to continually grow interest and development among packaged-application ISVs that have not yet committed to Itanium and IA-64, an issue with which Intel continues to grapple in 1999. Indeed, this will be a key determinant of whether Itanium is the harbinger to IA-64 success, or a wild success itself. For wherever enterprise applications are available on an enterprise platform, IS executives will follow. The key is for

forward-looking IS executives to begin assessing their IA-64 product plans today by working with their key system and application vendors.

To Aberdeen, Itanium heralds a new chapter in enterprise computing. In a nutshell, it will change the way that microprocessors, memory, compilers, and applications interact to support business-critical environments.

Appendix

IA-64 Architectural Overview

Explicitly Parallel Instruction Computing (EPIC) relies heavily on CPU-savvy compilers to optimize performance. The need for tight integration with compilers means that Intel is making an unprecedented effort to line up key industry players (system suppliers, ISVs, and services providers) so that Itanium and future IA-64 processors can provide comprehensive performance solutions out of the "starting gate."

Itanium also achieves a new level of parallel processing. New parallel-processing features include hardware support for multiple parallel data streams; fewer memory latencies than traditional architectures; better anticipation of instruction streams that can be made parallel; increasing the number of parallel streams supported; a third level of caching; and making data parallel before it enters the microprocessor.

Aberdeen concludes that Itanium and future IA-64 microprocessors will change the way IS executives look at Intel Architecture platforms. For example, most IS executives have used frequency speed (i.e., megahertz) on the Intel Architecture as one method of determining performance. With the introduction of IA-64, IS executives should consider measuring total system performance based on a number of factors that include other elements such as server I/O, memory, the instruction set itself, and key application performance.

IA-64's Key Architecture Components: Why Itanium Is a First

As a result of the introduction of IA-64, Aberdeen expects that frequency (MHz) will no longer be the primary decision point in the choice of CPU architecture. Instead, instructions per clock cycle (IPC) will become a more important measure of system efficiency and power as calculated with the frequency.

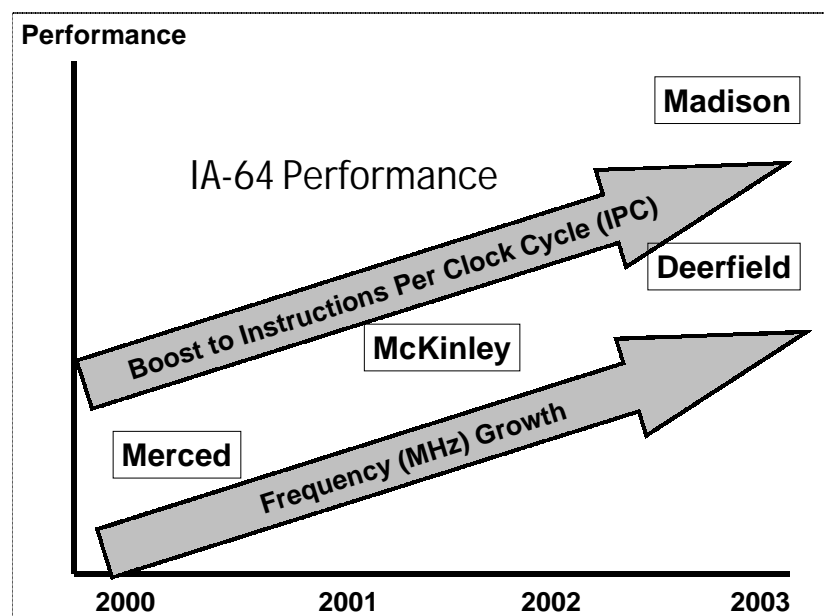
A number of key attributes make the IA-64 architecture different from most other microprocessor architectures now available to enterprise IS executives. First, IA-64 is not a pure RISC one-instruction-per-clock-cycle processor; so, to arrive at a quick-and-dirty estimate of performance, users should focus on IPC and clock frequency (the traditional rule of thumb). In the second place, the IA-64 hardware is simpler — because the compilers simplify and optimize the instructions specifically for IA-64 before the CPU processes them.

EPIC enables Itanium and future IA-64 microprocessors to handle more IPCs (that is, more application code) than previous microprocessors, and to feed these instructions to multiple on-chip functional units for execution on every clock cycle. Intel's IA-64 architecture designers are building significant performance

headroom into the architecture, allowing them to cram in more IPCs over the next few generations of IA-64 (Figure A).

In traditional architectures, a key limit to performance and scalability has been that the architecture feeds instructions and data to the CPU too slowly to fully stretch the ability to make data parallel. As a result, in most of today's microprocessor architectures, the CPU has the additional task of parallelizing instructions on the fly — a task that consumes valuable real estate and processing power on the CPU.

Figure A: IA-64 Performance Focused on Instructions Per Clock Cycle



Source: Aberdeen Group, July 1999

Much of the organizing that today's RISC and CISC processors do is called "pipelining." Pipelining examines a set of instructions, determines which parts of it can be performed simultaneously, and feeds those parts down different "pipelines" (parts of the CPU), each of which carries out the instructions in one part. Note that when the instruction set is fully parallelizable, all of the pipelines are working simultaneously, and the CPU's processing speed is multiplied by the number of pipelines. Note also that the CPU can use pipelining when the instruction stream comes to a "decision point" where it can take one of two possible paths: one pipeline will handle one path, and the other pipeline will handle the other.

By building more and more pipelines, current RISC and CISC architectures have been able to continually boost performance. But these architectures have not been able to fully parallelize most instruction streams — leaving a large potential performance boost on the table.

IA-64's Crown Jewel: Better Efficiency in Processing Application Code

With IA-64, the compiler takes the bulk of the responsibility for parallelizing application code. The compiler's optimization phase "pre-parallelizes" application code *specifically for the IA-64 architecture* before it is sent to the CPU. In effect, the compiler gives the CPU the blueprint for parallelizing the instruction stream *in the instruction stream itself*, and the CPU carries out that blueprint, adding its own optimizations. This process is particularly effective where the big opportunities for performance speed up are in parallelizing large chunks of code, because CPUs cannot "look ahead" very far in the instruction stream to detect these code chunks.

It should take some time for the new compiler-driven optimization to translate into major jumps in parallelization and performance. History has shown that compilers' programmers take two years (or more) to learn how to optimize effectively. Since the compiler suppliers — typically the hardware suppliers and Microsoft — should be about 18 months into the learning process when Itanium arrives, expect continuing performance improvements whenever a new compiler arrives, not just when a new IA-64 chip or upgrade ships.

Intel has introduced the following key technologies to boost the number of instructions that the processor can handle, and Intel has introduced a number of other key technologies that build on EPIC. These technologies include:

- Predication;
- Speculation;
- Rotating registers and other processing efficiencies; and
- Hardware enhancements (larger caches, enhanced scalability, larger integer and floating point units).

Predication

As noted in the previous section, a key factor in the ability of the CPU to maximize instruction parallelism is making the right decision about how to handle decision points. Traditional architectures (especially today's RISC ones) cannot see far ahead in the instruction stream, and therefore must guess much of the time which instruction stream after the decision point is the correct one. To avoid predicting incorrectly which branch/instruction stream to execute, today's CPUs may attempt to guess more intelligently ("branch prediction"), run both instruction streams in different pipelines ("speculative execution"), or execute one or both of the branches when the CPU is underutilized ("out-of-order execution"). Much of the time, these tactics do not fully leverage CPU parallelism: Intel estimates that these architectures' "mispredicted branches" (executing the wrong instruction stream) can cost up to 40% in CPU performance.

Predication takes a new approach to the problem — “looking ahead” in the compiler before the CPU receives the code. Predication can remove branches wherever possible. IA-64 can also tell the CPU more about each branch so that the CPU can do a better job of guessing the right instruction stream. For example, in a 100-interaction loop, the compiler would optimize for the 99 successes before branching on the 100th.

Moreover, Itanium's larger real estate means that it can simultaneously operate more pipelines than can most of today's traditional microprocessors. This feature gives the CPU the option, when all else fails, of doing speculative execution on a grander scale — squeezing out more parallelism when branches are particularly hard to predict.

Aberdeen finds that applications involving complex, frequently branching code — legacy applications, technical/scientific ones, and computation-bound solutions like economic order-quantity computation in ERP, stock-market analysis, and data analysis for data mining — will particularly benefit from predication.

Speculation

Instructions and data that arrive at the CPU late (from main memory or cache) typically slow performance. For example, the bus that brings the next instruction to the CPU may work best when each instruction to be sent to the CPU is stored right after the previous instruction; a branch will delay transmission of the next instruction until the bus finds the right instruction location. Or, the bus may be slower than the CPU; it cannot send instructions fast enough to keep the CPU busy all the time.

Itanium's speculation technology loads instructions from main memory or cache before the CPU is ready for them — even if it turns out that the CPU will not need to execute the instruction. In effect, speculation caches instructions and data on the CPU itself. Note that Itanium also can combine predication and speculation, eliminating many branches through predication, then speeding speculative execution or improving branch prediction for the remaining branches through speculation.

Speculation is particularly effective with applications that access cache frequently, because they need to quickly feed instructions and data to the CPU. For example, large-scale Internet or database applications must switch rapidly between large numbers of end-users; speculation can “pre-send” the instructions and data for each new end-user to the CPU so that the system can react more rapidly.

Rotating Registers and Other Processing Efficiencies

Itanium supports 128 64-bit-wide registers that system programmers can access. These registers allow the CPU to process more data at once. IA-64 designers have

added a number of monitoring and management features that will allow programmers to use these registers more effectively. For example, Itanium provides "rotating registers" that act like an assembly line, allowing the CPU to apply an operation to each register in turn. This approach means that when the same set of operations is to be applied to a stream of data, the CPU can "overlap" the operations for a high degree of parallelism. Key application functions such as large-database sorts and technical/scientific floating-point calculations (a necessity for many of today's workstation applications) will benefit from this feature.

Itanium also includes a number of performance-monitoring registers. For example, certain key registers keep a count of instructions per minute and cache misses. It is likely that in the future systems management software will be able to monitor these registers in real time, avoiding systems-management overhead and improving a user's ability to "tune" CPU performance to the workload.

Hardware Enhancements

Itanium and IA-64 include a number of hardware enhancements, including scalability features that Intel's OEM suppliers will likely use effectively. A number of system suppliers are building servers that will scale well beyond Intel's Standard High Volume (SHV) platforms of four- and eight-way Symmetrical Multiprocessor (SMP) architectures. To support this scalability, Itanium includes a number of memory and bus architecture enhancements. For technical/scientific computing environments, Itanium provides higher precision floating-point and integer-point units that can more effectively handle the large/precise-number computations typical of these environments.

Itanium has been designed to support better memory utilization — a crucial feature as system designers scale servers with up to 512 CPUs. This memory utilization capability includes the use of a 64-bit flat address space to handle larger applications and memory expansion. Aberdeen anticipates that Itanium will support flexible page sizes up to 256 MB to reduce paging overhead. And architects introduced a third-level cache to reduce the system-bus traffic that has limited a number of traditional RISC architectures. In fact, Level 3 Cache is a major cause of Itanium's immediate performance benefits.

Second, Itanium (and the IA-64 architecture) includes a highly efficient bus. Part of the specification lays out a dedicated, full-CPU-speed Level 2 bus that is meant solely for SMP traffic between processors. Furthermore, bus transactions are more effective because there is deferred-transaction support on the processor.

Finally, IA-64's Itanium processor will support 128 floating-point registers in addition to 128 integer registers, resulting in more write and read ports. For applications such as technical-workstation solutions and Internet-focused Digital Content Creation applications, the graphics quality, degree of rendering accuracy, and

speed of processing will be enhanced. For the larger-data-size applications typical of technical and scientific markets, Itanium will support two extended-precision-floating-point multiply and accumulate instructions, which are capable of single, double, and extended precision-floating-point calculations.

Itanium **Intelligent Hardware Reliability Features**

The processor, firmware, and the IA-64-enabled operating system work together to contain and fix errors in the IA-64 architecture. The Itanium processor signals the error to the firmware, which can then either correct the error or supply information to the OS for software correction. This feature provides additional opportunities to correct errors instead of crashing, resulting in greater system availability.

Another reliability enhancement, called data poisoning, is the flagging of double-bit ECC errors in the processor caches. Data poisoning allows the system to continue without a reboot; only the offending process needs to be terminated via machine check abort (MCA) when the poisoned data is consumed (i.e., when the poisoned data is accessed from the caches). A special processor state allows unaffected processes to continue, instead of bringing down the entire system. Other features include error logging — giving IS administrators the ability to trace application, I/O device, OS, and transaction errors — so that they can determine where faults are occurring in the system.

And, while IA-64 designers have focused on processor-centric enhancements, other Intel and OEM hardware system designers have started to explore ways to eliminate reliability issues found in today's server I/O architectures. Existing industry initiatives, including Next-Generation I/O and Future I/O, are working to introduce switched-fabric I/O-system backbones to replace server-PCI bus architectures over time. It is likely that shortly after the arrival of IA-64, systems that take advantage of these new I/O architectures will be delivered to market. System suppliers will also offer a number of other reliability features that have become common for enterprise servers, such as redundant fans, PCI Hot Plug components, redundant power supplies, and other reliability features.

*Aberdeen Group, Inc.
One Boston Place
Boston, Massachusetts
02108
USA*

*Telephone: 617 723 7890
Fax: 617 723 7897
www.aberdeen.com*

*© 2000 Aberdeen Group, Inc.
All rights reserved
March 2000*

Aberdeen Group is a computer and communications research and consulting organization closely monitoring enterprise-user needs, technological changes and market developments.

Based on a comprehensive analytical framework, Aberdeen provides fresh insights into the future of computing and networking and the implications for users and the industry.

Aberdeen Group performs specific projects for a select group of domestic and international clients requiring strategic and tactical advice and hard answers on how to manage computer and communications technology.